

What is claimed is:

1
2 1. A method of shutting down a node in a clustered computer system, the
method comprising:

3 (a) detecting a failure in a first node among a plurality of nodes in a
4 clustered computer system with a first group member resident on the first
5 node;

6 (b) in response to detecting the failure, transmitting a signal to each of
7 the other nodes in the plurality of nodes to initiate on each of the other nodes a
8 node leave operation that terminates clustering with the first node; and

9 (c) in response to detecting the failure, preemptively terminating a
10 second group member resident on the first node prior to any detection of the
11 failure by the second group member.

1 2. The method of claim 2, wherein the first group member is a cluster control
2 group member.

1 3. The method of claim 2, wherein the first group member is a member of a
2 group other than a cluster control group, wherein the method further comprises
3 sending a shutdown message to a cluster control group member resident on the first
4 node, and wherein transmitting the signal to each of the other nodes and preemptively
5 terminating the second group member are each initiated by the cluster control group
6 member in response to the shutdown message from the first group member.

1 4. The method of claim 1, wherein detecting the failure comprises detecting
2 an error that requires clustering to terminate on the first node.

1 5. The method of claim 4, wherein the error is selected from the group
2 consisting of a communication error, a protocol error, failure or termination of a
3 cluster name server, failure or termination of a cluster control group member, failure
4 or termination of a cluster liveness monitor, corruption of a critical cluster object,
5 and combinations thereof.

6 6. The method of claim 1, wherein transmitting the signal comprises
7 transmitting a distress signal to each other node in the plurality of nodes.

1 7. The method of claim 1, further comprising, at each of the other nodes in the
2 plurality of nodes, processing the distress signal on such other node by initiating a
3 node leave operation for each group member resident on such other node that is a
4 member of a group that has another member resident on the first node.

1 8. The method of claim 1, further comprising, at another node in the plurality
2 of nodes, performing a dependent failover in response to receiving the signal from the
3 first node.

1 9. The method of claim 1, wherein preemptively terminating the group
2 member comprises notifying each active group member resident on the first node that
3 clustering is ending on the first node.

1 10. The method of claim 9, further comprising ending clustering on the first
2 node after notifying each active group member resident on the first node.

1 11. The method of claim 9, wherein notifying each active group member
2 includes sending an error message to each active group member.

1 12. The method of claim 1, further comprising disregarding a second failure
2 detected from a group member, other than the first group member, that is resident on
3 the first node if transmission of the signal has already been initiated.

1 13. A method of shutting down a node in a clustered computer system, the
2 method comprising:

- 3 (a) in a group member resident on a first node among a plurality of
4 nodes in a clustered computer system, initiating a shutdown of the first node;
5 and
6 (b) shutting down the first node in response to initiation of the
7 shutdown by the group member.

1 14. The method of claim 13, further comprising detecting a failure in the first
2 node with the group member, wherein initiating the shutdown of the first node is
3 performed in response to detecting the failure.

1 15. The method of claim 14, wherein shutting down the first node comprises:

- 2 (a) transmitting a signal to each of the other nodes in the plurality of
3 nodes to initiate on each of the other nodes a node leave operation that
4 terminates clustering with the first node; and
5 (b) preemptively terminating a second group member resident on the
6 first node prior to any detection of the failure by the second group member.

1 16. The method of claim 15, wherein initiating the shutdown of the first node
2 comprises sending a shutdown message from the group member to a cluster control
3 group member, and wherein shutting down the first node further comprises:

- 4 (a) notifying a clustering infrastructure resident on the first node using
5 the cluster control group member to transmit the signal and initiate preemptive
6 termination of the second group member; and
7 (b) terminating the clustering infrastructure.

1 17. The method of claim 15, wherein transmitting the signal comprises
2 transmitting a distress signal to each other node in the plurality of nodes, the method
3 further comprising, at each of the other nodes in the plurality of nodes, processing the
4 distress signal on such other node by initiating a node leave operation for each group

5 member resident on such other node that is a member of a group that has another
6 member resident on the first node.

1 18. The method of claim 15, wherein preemptively terminating the second
2 group member comprises notifying each active group member resident on the first
3 node that clustering is ending on the first node.

1 19. The method of claim 15, further comprising disregarding a second failure
2 detected from another group member resident on the first node if the signal has
3 already been transmitted.

1 20. An apparatus, comprising:

2 (a) a memory accessible by a first node among a plurality of nodes in a
3 clustered computer system; and

4 (b) first and second group members resident in the memory, the first
5 group member configured to detect a failure in the first node; and

6 (c) a program resident in the memory, the program configured to shut
7 down the first node in response to the detected failure by transmitting a signal
8 to each of the other nodes in the plurality of nodes to initiate on each of the
9 other nodes a node leave operation that terminates clustering with the first
10 node, and preemptively terminating the second group member resident on the
11 first node prior to any detection of the failure by the second group member.

1 21. The apparatus of claim 20, wherein the first group member is a member of
2 a group other than a cluster control group, wherein the first group member is
3 configured to send a shutdown message to a cluster control group member resident on
4 the first node, and wherein the cluster control group member is configured to initiate
5 transmission of the signal to each of the other nodes and termination of the second
6 group member in response to the shutdown message from the first group member.

1 22. The apparatus of claim 20, wherein the program is configured to transmit
2 the signal by transmitting a distress signal to each other node in the plurality of nodes.

1 23. The apparatus of claim 20, wherein the program is configured to terminate
2 the group member by notifying each active group member resident on the first node
3 that clustering is ending on the first node.

1 24. The apparatus of claim 23, wherein the program is further configured to
2 end clustering on the first node after notifying each active group member resident on
3 the first node.

- 1 25. The apparatus of claim 20, wherein the program is further configured to
- 2 disregard a second failure detected from a group member, other than the first group
- 3 member, that is resident on the first node if transmission of the signal has already been
- 4 initiated.

1 26. A clustered computer system, comprising first and second nodes coupled
2 to one another over a network, wherein:

3 (a) the first node is configured to shut down in response to a failure
4 detected in the first node by a first group member resident on the first node by
5 transmitting a signal to the second node and preemptively terminating a second
6 group member resident on the first node prior to any detection of the failure by
7 the second group member; and

8 (b) the second node is configured to initiate a node leave operation
9 that terminates clustering with the first node in response to the signal from the
10 first node.

1 27. The clustered computer system of claim 26, wherein the second node is
2 configured to initiate a node leave operation for each group member resident on the
3 second node that is a member of a group that has another member resident on the first
4 node.

1 28. The clustered computer system of claim 27, wherein the second node is
2 configured to perform a dependent failover in response to receiving the signal from
3 the first node.

1 29. A program product, comprising:

2 (a) first and second group members, the first group member
3 configured to detect a failure in a first node among a plurality of nodes in a
4 clustered computer system;

5 (b) a program configured to shut down the first node in response to the
6 detected failure by transmitting a signal to each of the other nodes in the
7 plurality of nodes to initiate on each of the other nodes a node leave operation
8 that terminates clustering with the first node, and preemptively terminating the
9 second group member resident on the first node prior to any detection of the
10 failure by the second group member; and

11 (c) a signal bearing medium bearing the program.

1 30. The program product of claim 29, wherein the signal bearing medium
2 includes at least one of a recordable medium and a transmission medium.